

The post [Distributed Pumping: The "Swiss Army Knife" of Liquid Cooling](#) appeared first on [HPCwire](#).

Intel Collaborates with CERN to Support Upgraded LHC Experiments

Fri, 11/04/2016 - 12:49

Much of the media attention given to the particle accelerator experiments that happen at the European Organization for Nuclear Research, known as CERN, is focused on the Large Hadron Collider (LHC). It's no surprise, given the LHC is the world's largest, most complex machine, unravelling some of the toughest scientific problems by accelerating particles (protons or heavy ions) and making them collide in a gigantic 27-kilometer ring. But the work that happens immediately after particles collide in the LHC is not only critical to science, it's also quite interesting and important from a computing and data processing perspective. After all, the creation of particles or results in the LHC is only significant if scientists can quickly isolate them from millions of inconsequential signals for further study. That means ongoing advancements in trigger and data acquisition systems are essential to fully reaping the rich potential of the LHC. And, as you can imagine, the networking and computing challenges are extreme in nearly every dimension.

Historically, CERN's trigger and data acquisition systems have relied heavily on custom technology. For the next upgrade cycle scheduled to start around the end of 2018, however, CERN engineers and scientists were aware that the scalability limitations and costs of their custom solutions could start to limit progress. This led to a new collaborative project with Intel, through [CERN openlab](#), focused on exploring the feasibility of complementing one of the LHC's detectors with off-the-shelf data acquisition, data movement, and data filtering technology from Intel, including an FPGA platform. If the proof of concept is successful, it could have a significant impact on the future design and efficiency of trigger systems, including the remaining LHC detectors and other scientific instruments.

The Large Hadron Collider tunnel is located 100 meters underground on the Franco-Swiss border, near Geneva. Source: CERN

Preparing for enormous data growth

Simply phrased, the LHC fires high-energy particle beams at each other in a 27-kilometer ring. The detectors on the LHC include tracking devices that plot the trajectory of particles following collisions, as well as calorimeters that measure their energy, which helps to narrow down their identity. Niko Neufeld, a deputy project leader at CERN who works on the LHCb experiment, [likens the computing challenges](#) for the LHC experiments to solving millions of small puzzles involving up to a billion proton-proton collisions every second to retain the most interesting ones for deeper analysis.

Given the task at hand, the near-detector "online" computing challenges at CERN have always been extreme. And when CERN upgrades the LHC and detectors from late 2018 to early 2021 as part of its regular upgrade cycle, the data rates running through the various systems will jump significantly. Neufeld provides perspective on the jump: "Network scaling needs have really grown. For example, the largest networks we currently run at CERN have total bandwidth of around 800 GB per second. Following the upgrade work, our networks will need to carry between 40 and 50 TB of data per second. If you compare that to a Google data center, it may not sound impressive, but for a scientific instrument it's a huge step in terms of bandwidth." Neufeld said that the computing challenges have also grown quite complex. "We cannot simply scale up computing by a factor of 100 because we have whatever Moore's law gives us... We may be able to grow our computing farms by a factor of 1 or 2, but not much more. The rest has to come from more clever processing models," he added.

The Compact Muon Solenoid (CMS) is one of two large general-purpose detectors on the LHC. The image above captures a candidate proton-proton event as a part of the CMS search for the Higgs boson. Source: CERN

According to Neufeld, the detector teams face a host of challenges in preparing for the data-rate jump. Neufeld said that the customized trigger (hardware and software) on the front-end of the existing detector system had a long and expensive refresh cycle. "The engineering resources for ASICs [application-specific integrated circuits] and the FPGAs [Field Programmable Gate Arrays] in high energy physics are limited compared to industry, and the tight integration with the detectors makes upgrades outside of our major maintenance periods impractical," Neufeld explained. "We thought that moving to a more software-centered approach using off-the-shelf technology could greatly reduce these limitations and expand the developer base that is available in our community. Physicists are usually knowledgeable in some programming language, however, HDL [hardware description language] is a different challenge with a very long and steep learning curve."

Olof Barring, a deputy group leader of computing facilities at CERN, added that cost and energy considerations were also an important part of the equation. The lab needs to continue addressing greater computing, data moving and storage challenges with a more or less flat budget and within the datacenter's existing energy envelope.

Exploring options for three critical challenges

In 2014, the former CTO of CERN openlab, Sverre Jarp, invited Intel's Karl Solchenbach, director Exascale Labs Europe, and Steve Pawlowski, former vice president of advanced computing solutions at Intel, to discuss CERN's technical challenges with near-detector online computing to see if any Intel technology might be useful in addressing them. During the meeting, Neufeld presented three main challenges, and the participants worked together to map technologies to the challenges.

Challenge one: real-time or near-real-time data-processing with very short (order of 10 microseconds) latencies

The Intel team considered an Intel Xeon processor and Altera Stratix V FPGA integration along with another spatial architecture, and ended up recommending the Intel Xeon processor/FPGA configuration. Neufeld says that the reasoning was that the Intel Xeon processor/FPGA configuration should allow LHC experiments to potentially replace part, and in some cases all, of the custom electronics used in the first step of online data filtering. "That meant we would be able to use off-the-shelf hardware programmed using 'high-level' general purpose languages instead of HDLs, which was an important step for us," he explained.

Challenge two: very high throughput local area networks

For data transport, Intel recommended investigating the potential of Intel Omni-Path Architecture (Intel OPA) as an alternative to deep-buffer Ethernet switching. Neufeld said the main driver for considering Intel OPA over traditional approaches was cost. "It is certainly technically possible to build the network in a more traditional way, but it has become prohibitively expensive, given our budget," he explained.

Challenge three: the need of massive data-processing for data-reduction

For the filtering of detector data in software, the new Intel Xeon Phi processor was an attractive potential solution, given that the filtering process itself is quite parallel and individual collisions in the LHC are statistically independent.

Not your average POC

Once the CERN and Intel teams agreed that the identified off-the-shelf solutions for each challenge had potential, the hard work of proving the viability of each solution in the next-generation data acquisition environment needed to begin in earnest. In 2015, CERN and Intel decided to expand their existing collaboration through CERN openlab and signed an agreement for a joint three-year project called the High Throughput Computing Collaboration (HTCC). As of writing, the HTCC project has reached the halfway point and the core team, which includes seven CERN scientists and one Intel engineer on site, has made progress in each key challenge area.

The main CERN datacenter. The 110,000 processor cores and 10,000 servers hosted in its three rooms run 24/7. Source: CERN

Networking milestones

In the networking area, Neufeld said a big part of the challenge is simply getting access to a system big enough to test software. "To properly prepare our software, we need access to complete supercomputers. Fortunately, Intel can provide access to clusters that are up to the job," he explained. Neufeld said that the team recently had an important success while running its software on an Intel OPA cluster with more than 500 nodes. "We have already been able to achieve a full-duplex, high-throughput transit of 70 TB per second of data flying through the cluster, so that's already half of what we need by the time of the upgrades."

FPGA-related milestones

When CERN updates the LHC and the experiments from late 2018 to early 2021, its detectors will support trigger-free readouts. The LHC generates up to around 1 billion collisions per second in the experiments, and the goal is to read them all. Moving forward, a flexible software-based trigger system running in a large (up to 4000 nodes) computing farm will select the interesting collisions. In the meantime, CERN is investigating which technology options are the best fit for accelerating its algorithms. For its initial FPGA proof of concept, the HTCC team deployed a two-socket Intel Xeon processor/FPGA machine, which included the following hardware connected by the Intel QuickPath Interconnect:

- Intel Xeon CPU E5-2680 v2
- Altera Stratix V GX A7 FPGA with 234,720 adaptive logic modules (ALMs)

The team is particularly interested in the potential compute and power efficiency gains that are possible with using OpenCL in a combined CPU and FPGA system. With respect to the FPGA platform testing, Neufeld said that the team dealt with host of technical challenges because it wanted to do a meaningful comparison of OpenCL and to get a sense of the costs using a high-level framework. "Just for illustrative purposes, it took us two weeks to set up a new kernel using OpenCL compared to three months to complete an equivalent Verilog implementation, and we had a very skilled engineer on that job," explained Neufeld. "In the end, I think it was a good investment because we needed to prove to the electronic engineers that the new technology actually provided a less painful way to get the results we need."

Cherenkov angle reconstruction is used for particle identification in the detectors. Based on initial tests the Xeon/FPGA machine shows promise for processing greater numbers of photons after the LHC upgrade. Source: CERN

Following a series of test cases, such as sorting and calculating the Mandelbrot fractal, to understand the potential of the Xeon/FPGA system, the HTCC team developed an FPGA fine-tuned for the rigors of RICH (ring-imaging Cherenkov) reconstruction. Only then did it begin doing LHC-specific workload analysis.

In coming months, the HTCC team will also be testing a newer system that is built using a combined Intel Xeon CPU and FPGA in a single package. It will include the new high-performance Arria 10 FPGA from Altera as well as a faster interconnect of the CPU and FPGA.

Intel Xeon Phi processor milestones

Neufeld said that testing of the Intel Xeon Phi processor platform will be a major focus of the HTCC team for the next year or so. He noted that like everyone else, the team needs to figure out how to adapt well to the new architecture and different level of parallelism. To achieve this the HTCC team has been working with Intel engineers on benchmarking and understanding the different algorithms' implementations using Intel analysis tools, such as Intel VTune Amplifier XE and Intel Advisor XE, as well as different performance models, such as the roofline model.

"In addition to the inclusion of bootable sockets, in-package memory and high main-memory bandwidth, what is particularly interesting with Intel Xeon Phi processors is the integrated fabric and its potential to quickly distribute workloads to where they fit best. We will test both data movement aspects on the Intel Xeon Phi processor as well as the distribution of the algorithms between the Intel Xeon and Intel Xeon Phi processor using the fabric as an interconnect," Neufeld added.

Pushing the boundaries of precision

When asked what the progress made on the data acquisition systems for the LHC might mean for wider applications, Neufeld said the value is all about greater precision for complex experiments. "To some extent, it's just statistics. Either you really increase the amount of data and the precision by a significant factor, or you stop doing it," Neufeld explained. "This work should lead to an important jump in precision. For example, the LHCb experiment's online collection and analysis system currently selects just 1 million of the 40 million bunches of protons that cross in the accelerator every second, with the others being discarded based on less-precise hardware-calculated signatures. After the LHCb upgrade, the number of collisions is set to grow yet further, and we will look at all of them in the software, in order to take the best physics out of there. And since there are actually only a couple of milliseconds to do that for each collision, it's really quite a leap forward." Direct discussions with the Intel development team will continue to be invaluable to getting things right on projects as the team races to meet its deadline for the start of the upgrade work.

The post [Intel Collaborates with CERN to Support Upgraded LHC Experiments](#) appeared first on [HPCwire](#).

XSEDE Announces PEARC17 Conference

Fri, 11/04/2016 - 12:19

Nov. 4 – The Extreme Science and Engineering Discovery Environment (XSEDE) annual conference, most recently held as XSEDE16 in Miami last July, is transforming into an independent entity designed to unite the high-performance computing and advanced digital research community. The [Practice & Experience in Advanced Research Computing conference](#) (PEARC) will welcome all who care about using advanced digital services for research. The PEARC17 Conference will take place in New Orleans, Louisiana, July 9-13, 2017.

The goal in spinning off the PEARC (pronounced "perk") conference is to form a hub to address the many challenges in the field of advanced research computing. The operation and use of campus and national advanced research computing resources and services affects vastly different communities in a variety of ways, including: directors and managers, system administrators, user support staff, and facilitators, computational scientists, government agencies, students, and educators. To support this, PEARC participants will be comprised of academic campus, national and international cyberinfrastructure and research computing enablers to encompass a broad set of communities.

PEARC's inaugural year will feature support from a number of community organizations and will continue building on prior conferences' success and core audiences. Organizations who are supporting the new conference include the Advancing Research Computing on Campuses: Best Practices Workshop (ARCC), XSEDE, the [Science Gateways Community Institute](#), the [Campus Research Computing \(CaRC\) Consortium](#), the [ACI-REF](#) consortium, the [Blue Waters Project](#), [ESnet](#), [Open Science Grid](#), [Compute Canada](#), the [EGI Foundation](#), the [Coalition for Advanced Scientific Computation](#) (CASC), and [Internet2](#). Organizations interested in joining the PEARC community can submit their interest here.

In perhaps the most visible sign of support, ACI-REF and the National Center for Supercomputing Applications (NCSA) are co-locating the 2017 ARCC best practices workshop with PEARC17. Last year's XSEDE16 conference attracted more than 525 attendees from 44 states, and the 2016 ARCC workshop drew more than 100 attendees from universities across the country. PEARC17 will continue to support the community's interest in co-locating smaller meetings in this collective effort to build a self-sustaining, independent conference for the advanced research computing community.

"PEARC will create a forum owned by the community to foster exchanges around the 'state of the practice' in advanced research computing—discussing challenges, opportunities, and solutions," said John Towns, principal investigator for XSEDE. "Those who have attended the impactful annual XSEDE conference over the years will benefit and benefit from this new collaboration. We are looking forward to welcoming everyone at PEARC17."

The merging of these activities creates a unified forum for developing, delivering, and supporting infrastructure enabling science and will focus on best practices and the experiences that shape the technical aspects of research. Bringing all of these efforts into one event will allow it to expand beyond the scope of the XSEDE program and the National Science Foundation where XSEDE will continue to be a partner in facilitating but no longer own the conference.

"On behalf of the members of the ACI-REF consortium, we strongly support this effort to bring the broader cyberinfrastructure community together under one roof for an exchange of ideas and new discovery," said Gwen Jacobs, PhD, Director of Cyberinfrastructure, University of Hawaii. "In ACI-REF we have seen the benefit first hand of the gains made when groups come together to exchange expertise and learn from each other. Co-locating our ARCC meeting with PEARC17 will broaden the scope and impact of the meeting for all participants. We're excited to be part of this community effort."

"PEARC will allow for sharing the best practices as well as accomplishments and impacts, said Bill Kramer, project director and principal Investigator of the Blue Waters Project at the NCSA. "Furthermore, PEARC can be an important component in supporting the goals of the National Strategic Computing Initiatives, in particular the objectives for increasing the coherence between technology for modeling/simulation and data analytics, increasing the capacity and capability of an enduring national HPC ecosystem, and developing and/or expanding the U.S. government, industry, and academic collaborations to share the benefits of high performance research."

"Expanding the XSEDE conference to include the broader aspects of research computing in today's data-driven scientific environments promises to fill a critical need faced by scientists across the country," said Eli Dart, Network Engineer at ESNet Science Engagement Group. "Bringing the community together to share best practices benefits us all, from resource providers to facilitators and engagement engineers, to the scientists we serve."

"Working together with XSEDE and other partners, the Campus Research Computing Consortium (CaRC) hopes to help transform PEARC into a broader and more inclusive meeting that highlights the successes and interplay among campus, regional, national, and international CI," said Thomas E. Cheatham, III, Director, Research Computing and Center for High Performance Computing, and Professor of Medicinal Chemistry, at the University of Utah.